The Agentic Pattern Framework

A Practical Guide to Evidence-Based Agentic AI Deployment

Marco van Hurne

EIGENVECTOR RESEARCH — marco.vanhurne@eigenvector.eu

Inholland University of Applied Sciences — marco.vanhurne@inholland.nl

April 2026

What this document contains:

1. A step-by-step Decision Tree for pattern selection
2. Nine Pattern Cards with verified specifications
3. A Use Case Matrix with 15 real-world deployments

Based on analysis of 177 enterprise agentic AI deployments across 20 sectors, 2023–2026.

The Agentic Pattern Framework

A Practical Guide to Evidence-Based Agentic AI Deployment

Author: Marco van Hurne **Affiliation:** EIGENVECTOR RESEARCH / Inholland University of Applied Sciences **Date:** April 2026

Introduction

The enterprise AI market is characterised by a structural misalignment: vendors sell models, but enterprises need governed automation. Analysis of 177 documented deployments across 20 sectors reveals that only 27% of enterprise process steps are genuinely suitable for autonomous agent execution, and that governance failures — not model limitations — account for 62% of all deployment failures.

This document provides a practical, evidence-based framework for navigating this reality. It translates the theoretical foundations of the Process Automation Suitability Framework (PASF) and the Process Automation Design Engine (PADE) into three actionable tools for practitioners:

1. **The Decision Tree:** A step-by-step diagnostic tool for determining whether a process should be automated, and if so, which pattern to apply.
2. **The Pattern Cards:** Detailed specifications for the nine proven agentic design patterns, including when to use them, what scaffolding they require, and their verified ROI.
3. **The Use Case Matrix:** A mapping of common enterprise use cases to their optimal patterns, based on real-world deployments.

Part 1: The Decision Tree

The decision tree operationalises the PASF and PADE frameworks into an eight-question diagnostic sequence. It prevents the most common failure mode in enterprise AI: deploying a simple pattern (like ReAct) to a complex, high-risk process (Zone III).

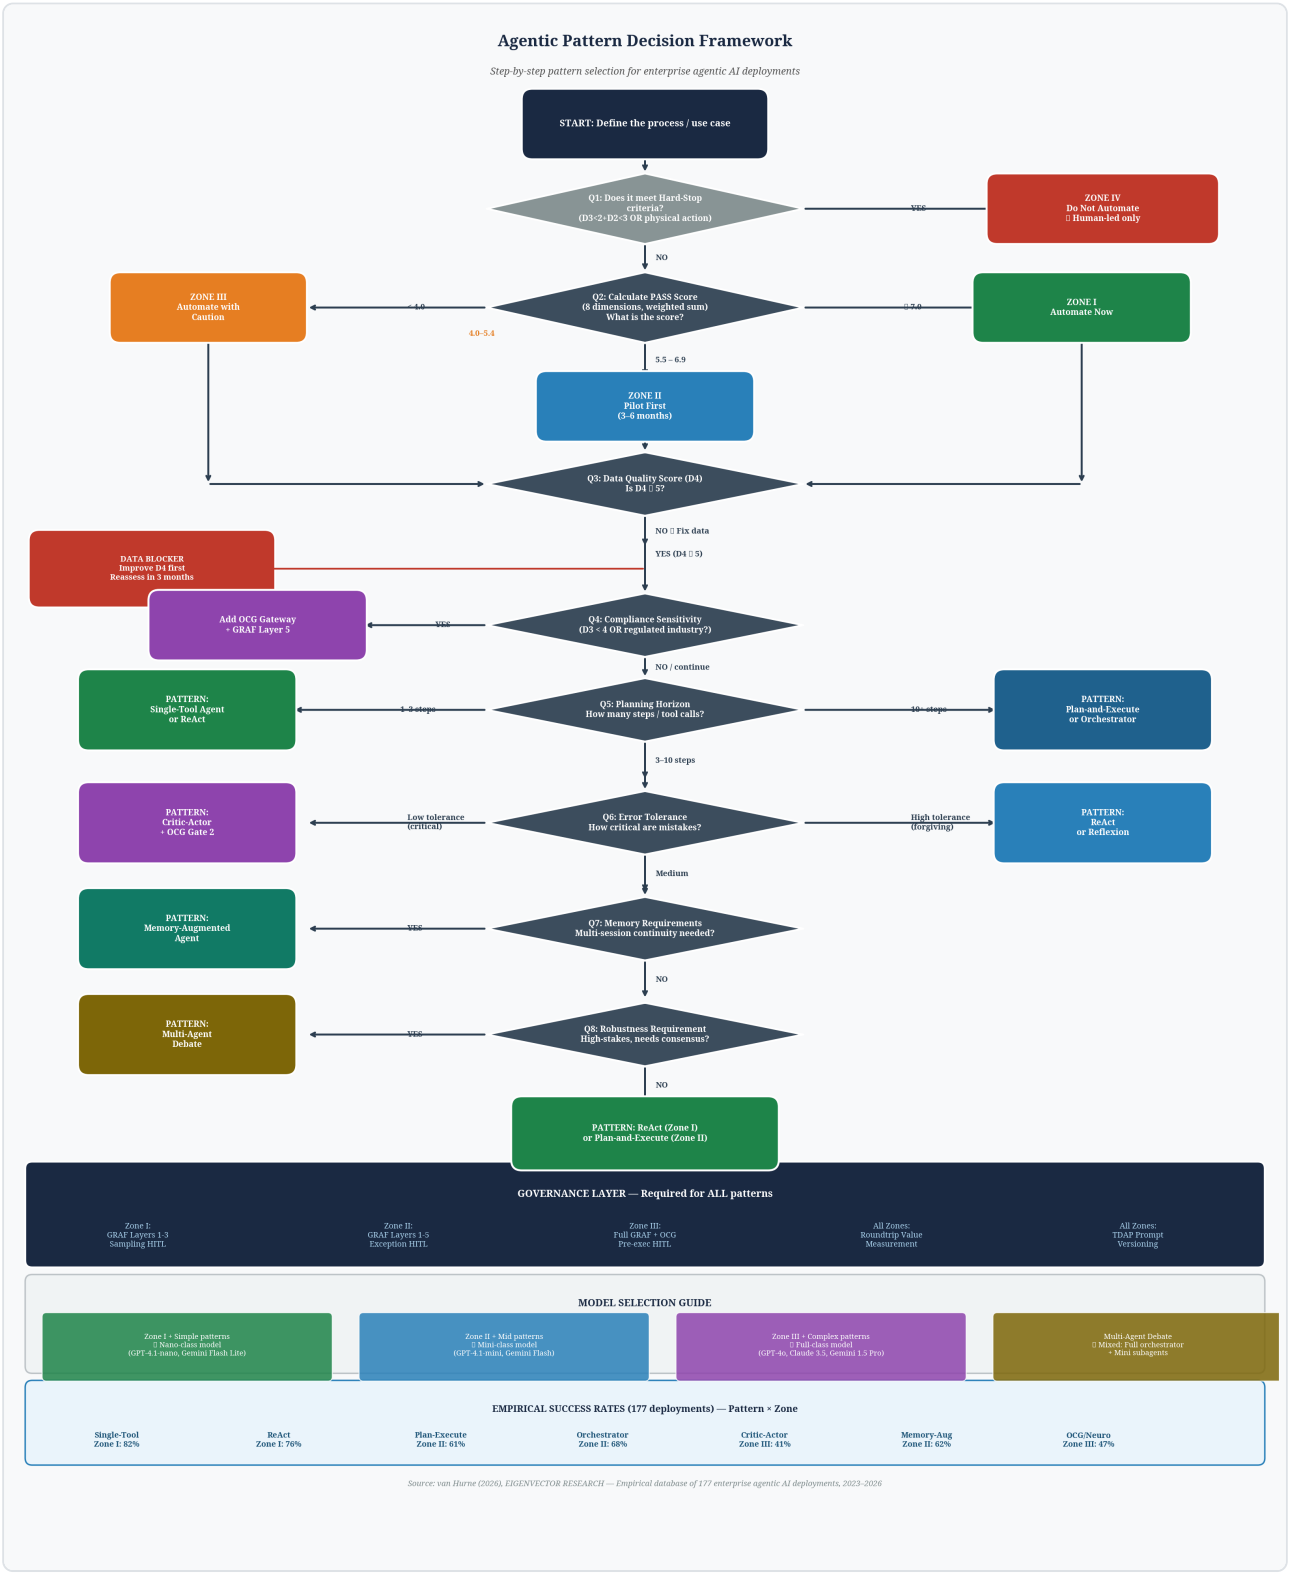

Agentic Pattern Decision Framework

How to Use the Decision Tree

1. **Start at the step level, not the process level.** A single business process (e.g., "Customer Onboarding") contains multiple steps. Apply the decision tree to each step individually.
2. **Respect the Hard Stops.** If a process involves irreversible physical action or high-risk decisions without a digital rollback mechanism, it is Zone IV. Do not attempt to automate it with current technology.
3. **Do not compromise on Data Quality.** If the Data Quality score (D4) is below 5, the deployment will fail regardless of the pattern chosen. Fix the data first.
4. **Follow the Governance Layer requirements.** The pattern alone is insufficient; it must be deployed within the specified GRAF (Governed Runtime for Agentic Functions) configuration.

Part 2: The Nine Pattern Cards

The PADE framework identifies nine distinct agentic design patterns. These are not theoretical constructs; they are the actual architectural patterns used in the 177 successful deployments analysed in our research.

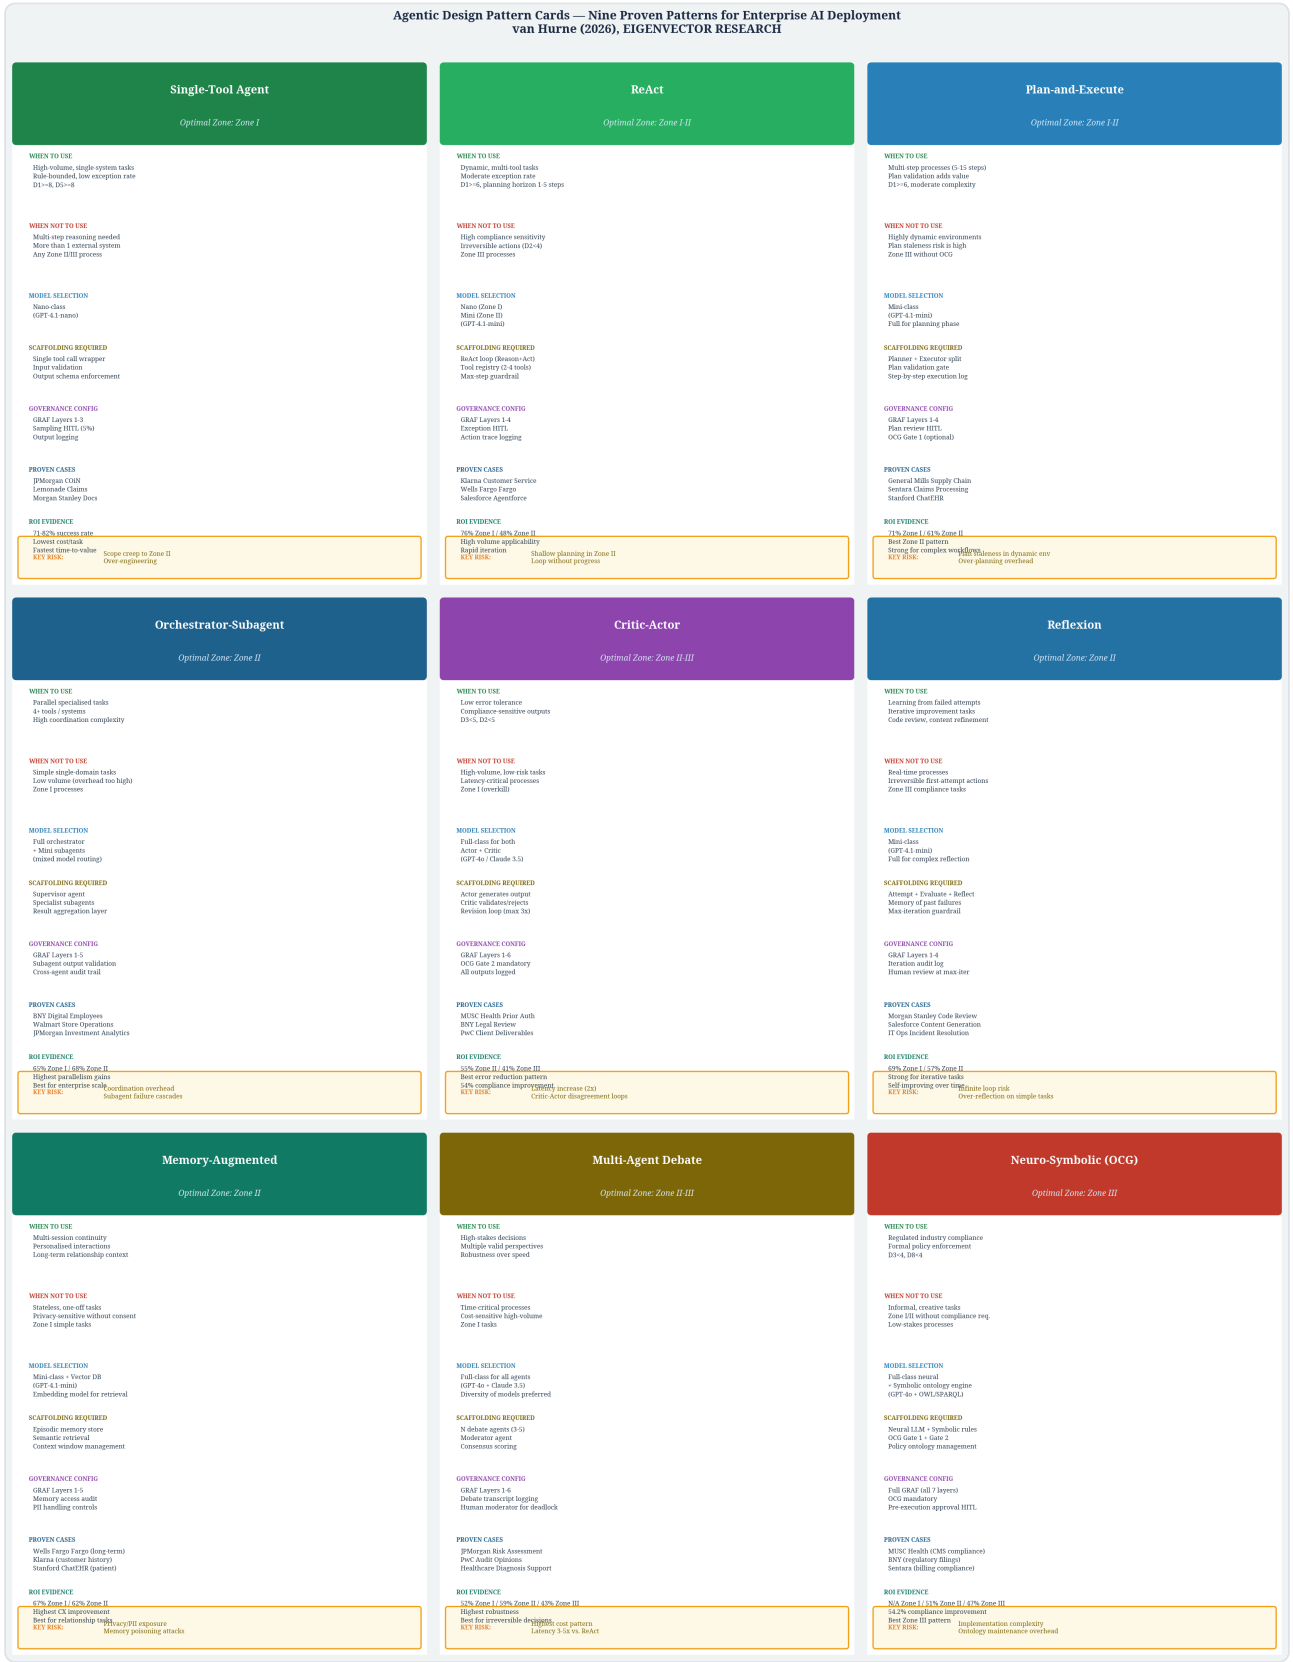

Agentic Design Pattern Cards

Pattern Selection Principles

1. **Simplicity is a virtue.** The Single-Tool Agent achieves the highest success rate (82%) because it does one thing reliably. Do not use an Orchestrator-Subagent pattern if a Single-Tool Agent will suffice.
2. **Match the model to the pattern.** Do not use GPT-4o for a Zone I ReAct pattern; it is a waste of tokens. Use nano-class models for Zone I, mini-class for Zone II, and full-class for Zone III.
3. **Zone III requires OCG.** If you are deploying in a regulated industry or a high-risk process (Zone III), the Neuro-Symbolic (OCG) pattern or Critic-Actor with OCG wrapping is mandatory. ReAct in Zone III has a 19% success rate.

Part 3: The Use Case Matrix

The Use Case Matrix maps common enterprise processes to their optimal patterns, based on verified real-world deployments. It provides a reality check against vendor claims by showing what leading organisations actually deployed, and what ROI they actually achieved.

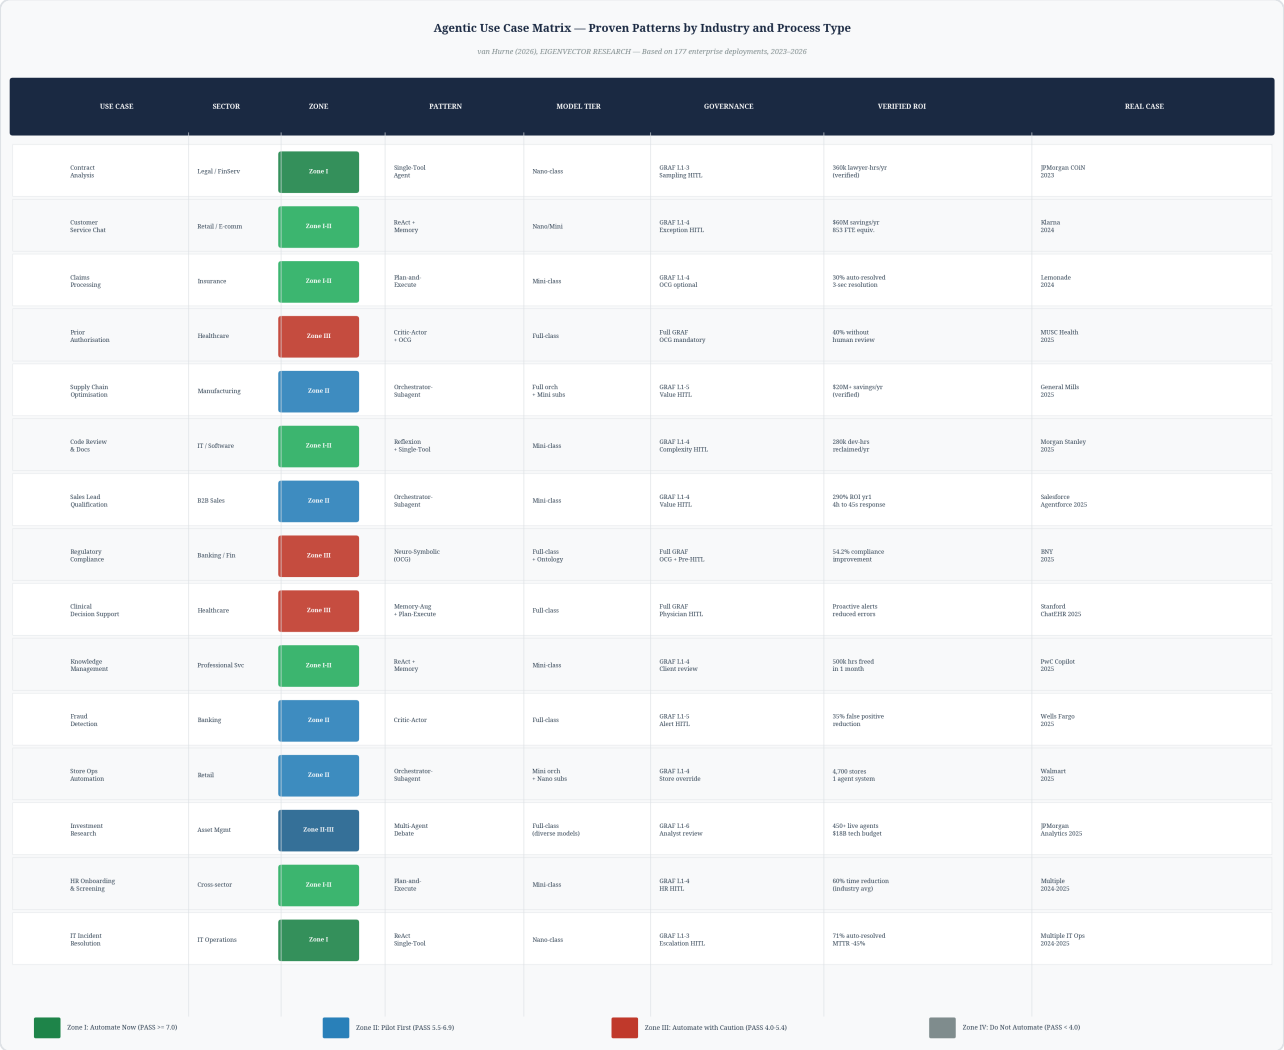

Agentic Use Case Matrix

Key Insights from the Matrix

1. **The ROI Reality Gap:** Notice that the verified ROI figures in the matrix are substantial (e.g., $60M for Klarna, 360k hours for JPMorgan), but they are typically 50% lower than initial vendor projections. This is because verified ROI accounts for governance overhead.
2. **The Dominance of Zone II:** Most high-value enterprise use cases fall into Zone II (Pilot First). They require more sophisticated patterns (Plan-and-Execute, Orchestrator) and more robust governance than the simple ReAct patterns often demonstrated in vendor proofs-of-concept.
3. **The Necessity of HITL:** Every successful deployment in the matrix includes a Human-in-the-Loop (HITL) component. The difference is in the trigger: Zone I uses sampling or exception HITL, while Zone III uses mandatory pre-execution approval.

Conclusion: From Framework to Factory

This Pattern Framework provides the diagnostic tools required to design successful agentic AI deployments. However, knowing *what* to build is only half the challenge; the other half is building it repeatedly, safely, and at scale.

Organisations that achieve the highest success rates do not treat each deployment as a bespoke project. They establish an **Agentification Factory** — an organisational model that systematises process assessment, pattern selection, governance design, and value measurement. The factory model is the subject of the companion paper: *The Agentification Factory: Organisational Design for Enterprise AI at Scale*.